Rule DAS300: Perhaps shared DASD caused performance problems

Finding: CPExpert believes that accessing conflicts caused by sharing DASD between systems or MVS images may have caused performance problems.

Impact: This finding is used to assess whether sharing DASD between systems or MVS images caused performance problems.

Logic flow: The following rules cause this rule to be invoked:

DAS100: Volume with the worst overall performance
DAS110: Seeking was major cause of I/O response delay
DAS120: Missed RPS reconnect was major cause of I/O delay
DAS130: Large PEND time was major cause of I/O delay
DAS150: Missed cache read hits was major cause of I/O delay

Discussion: DASD volumes can be shared between systems or between MVS images. Sharing of the DASD volumes might be implemented to allow backup of data, to facilitate recovery or restart, to permit transfer of data from one system to another, etc.

In some situations, sharing DASD volumes has little impact on performance (for example, few I/O operations might be directed to the shared volumes from potentially conflicting systems).

In other situations, sharing DASD volumes can have a significant impact on the performance of the shared volumes, and consequently, on the performance of the applications accessing the shared volumes.

CPExpert can perform an analysis of conflicts between DASD shared between systems or MVS images. The analysis performed by CPExpert is not intended to identify an isolated performance problem. Rather, CPExpert attempts to identify those problems that **continually** cause shared DASD performance problems. Shared DASD analysis is invoked by specifying **%LET SHARED = Y;** in USOURCE(DASGUIDE). Shared DASD analysis is an option, because more processing is required to perform shared DASD analysis.

CPExpert performs the following processing if you have indicated that an analysis of potential conflicts between shared DASD should be performed:

- CPExpert determines whether the "worst" devices selected for detailed analysis are shared with another system. If so, CPExpert performs an analysis of potential conflicts caused by shared DASD.

- CPExpert identifies other systems that reference the "worst" device. This identification is accomplished by analyzing the SMF Type 74 data in the performance data base relating to all other systems. The SMF Type 74 data contain the VOLSER for each device referenced. CPExpert simply selects SMF Type 74 information for the systems that reference the VOLSER of the "worst" device. This information is retained for more detailed analysis about potential conflicts.

  There is a potential (but very unlikely) problem with this method of identifying devices shared between systems. Multiple systems in the performance data base could use the same VOLSER to identify different devices. This could happen if the devices were not shared between systems.

  For example, suppose that CPExpert had identified PAGE01 as the "worst" device. Several system in the performance data base could reference VOLSER PAGE01, but the devices with VOLSER PAGE01 could be unique to each system. CPExpert would assume that all references by other systems to PAGE01 applied to the "worst" device being analyzed. The references could apply to a totally different device, and the other systems might not even share DASD with the system being analyzed.

  If this should be a problem (that is, if the DASD Component reports shared DASD conflicts with systems that do not share the device being analyzed), simply ignore the analysis produced by CPExpert[1].

- Once CPExpert has identified all systems that reference the "worst" device, CPExpert analyzes the DASD I/O characteristics of these systems with respect to the "worst" device. As described earlier, the analysis makes a basic assumption that the I/O activity from the different systems is random among the systems (for example, the code assumes that the I/O activity of System B is independent from the I/O activity of System A).

CPExpert will produce Rule DAS300 to list statistics relating to potential conflicts, by system, by volume, and by RMF measurement interval. The following example shows sample output from Rule DAS300:

[1]We do not feel that this problem will be common. It is described only to alert you to a potential incorrect analysis. If any user encounters this problem and it becomes annoying, code can be implemented to allow users to identify specific systems that share DASD with the system being analyzed. At present, this option seems to add unnecessary complexity to the user options.


```
RULE DAS300:  PERHAPS SHARED DASD CONFLICTS CAUSED PERFORMANCE PROBLEMS

   Accessing conflicts caused by sharing VOLSER PPVOL1 between systems
   might have caused performance problems for the device during the
   measurement intervals shown below.  Conflicting systems had the
   indicated I/O rate, average CONN time per second, average DISC time
   per second, average PEND time per second, and average RESERVE time
   to the device.  Even moderate CONN, DISC, or RESERVE can cause delays
   to shared devices.
                ..          I/O   MAJOR    OTHER  -------OTHER SYSTEM DATA--------
   MEASUREMENT INTERVAL    RATE  PROBLEM   SYSTEM I/O RATE  CONN  DISC  PEND  RESV
    8:30- 8:45,22OCT2001 131.6  PEND TIME J80       37.3   0.042 0.003 0.065 0.000
                                          JF0      147.0   0.129 0.005 0.158 0.000
                                          JH0      368.8   0.372 0.036 0.799 0.000
                                          Z0       459.7   0.406 0.017 0.765 0.000
                                          TOTAL   1012.8   0.949 0.061 1.786 0.001
    8:45- 9:00,22OCT2001 108.5  PEND TIME J80       41.2   0.046 0.003 0.066 0.000
                                          JF0      195.1   0.169 0.006 0.226 0.000
                                          JH0      411.7   0.411 0.032 0.718 0.001
                                          Z0       498.9   0.432 0.015 0.795 0.001
                                          TOTAL   1147.0   1.058 0.056 1.805 0.002
```

Rule DAS300 shows, for each RMF measurement interval, the I/O rate and the major problem during the RMF interval, of the device being analyzed. The remaining data shows relevant information (I/O rate, CONN time, DISC time, PEND time, and RESERVE time) for the other systems that reference the device.

CPExpert summarizes the other system data, into a TOTAL row for each RMF interval. In some instances, the TOTAL per second time in a particular interval will be more than one second. In the case of DISC time or PEND time, this situation is caused by multiple I/O operations being delayed for the device. This commonly happens only with Parallel Access Volume (PAV) devices.

However, the above example shows that the CONN time is larger than one second per second! CONN time is normally thought to involve data transfer between the device and the host system, and this concept is convenient for most analysis. Clearly, a device cannot be active transferring data for more than one second per second.

The CONN time actually is a hardware construct that is measured at the channel subsystem level, and includes all hardware protocol between the channel, the director port, and the control unit. The hardware protocol is a very small amount of time. However, if there are many I/O operations, and the hardware protocol must take place for each I/O operation (and takes place at several points in the I/O operation), the total hardware protocol can become more significant.

In the above example, there was a total of 1147.0 I/O operations shown in the last total line (which is the total for “other systems” referencing the device). If the hardware protocol connect time were multiplied by 1147 I/O operations per second, it is easy to appreciate that the total hardware protocol connect time from the multiple systems would add appreciably to the connect time.

Thus, while the connect time for a particular device cannot exceed one second per second, once the hardware protocol connect time is added to the device connect time, the total can exceed one second connect time per second for a very active device.

Suggestion: You should use the information displayed by Rule DAS300 to assess the significance of the performance problems caused by shared DASD.

